HP XC Systems with HP Server Blades and Enclosures HowTo

Version 3.1 or Version 3.2

Published: April 2007
Edition: 9

© Copyright 2006, 2007 Hewlett-Packard Development Company, L.P.
Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.
AMD and AMD Opteron are trademarks or registered trademarks of Advanced Micro Devices, Inc.

Firefox is a registered trademark of Mozilla Foundation.

InfiniBand is a registered trademark and service mark of the InfiniBand Trade Association.

Intel, Itanium, and Xeon are trademarks or registered trademarks of Intel Corporation in the United States and other countries.

Internet Explorer is a registered trademark of Microsoft Corporation

Linux is a U.S. registered trademark of Linus Torvalds.

Quadrics and QsNet^{II} are registered trademarks of Quadrics, Ltd.

Red Hat and RPM are registered trademarks of Red Hat, Inc.

UNIX is a registered trademark of The Open Group.

Windows and Internet Explorer are registered trademarks of Microsoft Corporation.

Table of Contents

1 Overview....9
1.1 Minimum Requirements....9
1.2 Read the Documentation Before You Begin....9
1.3 Supported Server Blade Combinations....10
1.4 c-Class Server Blade Hardware Components....10
1.4.1 Supported HP ProLiant C-Class Server Blade Models....10
1.4.2 Enclosures and Onboard Administrators....11
1.4.3 iLO2 Console Management Device....11
1.4.4 Management Processor Console Management Device....12
1.4.5 Mezzanine Cards....12
1.4.6 Interconnect Modules....12

2 Task Summary and Checklist....13
2.1 Best Practice for System Configuration....13
2.2 Installation and Configuration Checklist....13

3 Cabling....15
3.1 Network Overview....15
3.2 Cabling for the Administration Network....15
3.3 Cabling for the Console Network....16
3.4 Cabling for the Interconnect Network....17
3.4.1 Configuring a Gigabit Ethernet Interconnect....17
3.4.2 Configuring an InfiniBand Interconnect....18
3.4.3 Configuring the Interconnect Network Over the Administration Network....19
3.5 Cabling for the External Network....19
3.5.1 Configuring the External Network: Option 1....19
3.5.2 Configuring the External Network: Option 2....20
3.5.3 Configuring the External Network: Option 3 - Non Gigabit Ethernet Interconnect Clusters....21
3.5.4 Creating VLANs....22

4 Installing HP XC System Software On the Head Node....23
4.1 Task 1: Refer to the Installation Guide....23
4.2 Task 2: Install HP XC System Software on the Head Node....23
4.2.1 Connect to the Onboard Administrator....23
4.2.2 Start the Installation....24

5 Discovering the Hardware Components....27
5.1 Task 1: Prepare for the System Configuration....27
5.1.1 Node Naming Differences....27
5.1.2 Head Node Naming....27
5.2 Task 2: Change the Default IP Address Base (Optional)....28
5.3 Task 3: Use the cluster_prep Command to Prepare the System....28
5.4 Task 4: Discover Switches....29
5.5 Task 5: Set the Onboard Administrator Password....30
5.6 Task 6: Discover Enclosures and Nodes....30

6 Making Node-Specific Settings....33
6.1 Making Settings on Non-Blade Servers....33
6.2 Making Settings on HP ProLiant Server Blades....33
6.3 Making Settings on HP Integrity Server Blades....36

7 Configuring the HP XC System....39
7.1 Task 1: Install Patches or RPM Updates....39
7.2 Task 2: Refer To the Installation Guide For System Configuration Tasks....39
7.2.1 Using Specific IP Addresses to Configure InfiniBand Interconnect Switch Controller Cards....39
7.2.2 Running the startsys Command With Specific Options To Start the System and Propagate the Golden Image....39
7.3 Task 3: Verify Success....40
7.4 You Are Done....40

8 Troubleshooting....41
8.1 One or More Ports Do not Communicate Properly on a Gigabit Ethernet Switch....41
8.2 lsadmin limrestart Command Fails....41

A Configuration Examples....43
A.1 Gigabit Ethernet Interconnect With Half-Height Server Blades....43
A.2 InfiniBand Interconnect With Full-Height Server Blades....43
A.3 InfiniBand Interconnect With Mixed Height Server Blades....44

Index....47

List of Figures

3-1 Administration Network Connections....16
3-2 Console Network Connections....17
3-3 Gigabit Ethernet Interconnect Connections....18
3-4 InfiniBand Interconnect Connections....18
3-5 External Network Connections: Full-Height Server Blades and NIC1 and NIC2 in Use....20
3-6 External Network Connections: Half-Height Server Blades and NIC1 and NIC2 in Use....21
3-7 External Network Connections: Half and Full-Height Server Blades and NIC1 in Use....22
A-1 Gigabit Ethernet Interconnect With Half-Height Server Blades....43
A-2 InfiniBand Interconnect With Full-Height Server Blades....44
A-3 InfiniBand Interconnect With Mixed Height Server Blades....45

List of Tables

1-1 Minimum Requirements....9
2-1 Installation and Configuration Checklist....13
4-1 Head Node Installation Instructions....23
4-2 Boot Command Line Options Based on Hardware Model....25
6-1 Boot Order for HP ProLiant Server Blades....34
6-2 Additional BIOS Setting for HP ProLiant BL685c Nodes....35
6-3 Adding a Boot Entry and Setting the Boot Order on HP Integrity Server Blades....37
7-1 InfiniBand Switch Controller Card Naming Conventions and IP Addresses....39

1 Overview

HP Server Blade c-Class servers (hereafter called server blades) are perfectly suited to form HP XC systems. Physical characteristics make it possible to have many tightly interconnected nodes while at the same time reducing cabling requirements. Typically, server blades are used as compute nodes but they can also function as the head node and service nodes. The hardware and network configuration on an HP XC system with HP server blades differs from that of a traditional HP XC system, and those differences are described in this document.

This HowTo contains essential information about network cabling, hardware preparation tasks, and software installation instructions that are specific to configuring HP server blades for HP XC. HP recommends that you read this entire document before beginning.

1.1 Minimum Requirements

Table 1-1 lists the minimum requirements to accomplish the tasks described in this HowTo.

Table 1-1 Minimum Requirements

Component	Minimum Requirement
Software Version	Distribution media for HP XC System Software Version 3.1 or Version 3.2 that is appropriate for your cluster platform architecture
Hardware Configuration	• A hardware configuration consisting of HP server blades to act as compute nodes and possibly as the head node and service nodes • At least one ProCurve 2800 series switch, which is required at the root • Optional ProCurve 2600 series switches • Gigabit Ethernet or InfiniBand® interconnect switches • A local PC or laptop computer that is running a recent version of Mozilla Firefox® or Internet Explorer®
Knowledge and Experience Level	• You must have previous experience with a Linux® operating system. • You must be familiar with HP server blades and enclosures and related components by reading the documentation that came with your model of HP server blade.
Documentation	• The most recent version of this HowTo • The installation, administration, and user guides for the following components: — HP (ProLiant or Integrity) C-Class Server Blades — HP BladeSystem c-Class Onboard Administrator — HP Server Blade c7000 Enclosure • *HP XC System Software Release Notes* • *HP XC Hardware Preparation Guide* • *HP XC System Software Installation Guide*

1.2 Read the Documentation Before You Begin

Before you begin, HP recommends that you read the related documentation listed in Table 1-1 to become familiar with the hardware components and overall system configuration process.

If you do not have the required documentation in your possession, see the following sources:

- The most current documentation for HP Server Blades, enclosures, and other server blade components is available at the following Web site:

  http://www.hp.com/go/bladesystem/documentation

- The most current edition of the Version 3.1 or Version 3.2 HP XC System Software Documentation Set is available at the following Web site:

  http://www.docs.hp.com/en/highperfcomp.html

- This HowTo is evolving, so go to http://www.docs.hp.com/en/highperfcomp.html and make sure you have the latest version of this HowTo because it might have been updated since you downloaded the version you are reading now.

1.3 Supported Server Blade Combinations

The HP XC System Software supports the following server blade hardware configurations:

- A hardware configuration composed entirely of HP server blades, that is, the head node, the service nodes, and all compute nodes are server blades.
- A hardware configuration can contain a mixture of Opteron and Xeon server blades
- A mixed hardware configuration of HP server blades and non-blade servers where:
  - The head node can be either a server blade or a non-blade server
  - Service nodes can be either server blades or non-blade servers
  - All compute nodes are server blades

1.4 c-Class Server Blade Hardware Components

This section describes the various server blade components in an HP XC hardware configuration.

1.4.1 Supported HP ProLiant C-Class Server Blade Models

HP ProLiant C-Class server blades offer an entirely modular computing system with separate computing and physical I/O modules that are connected and shared through a common chassis, called an enclosure. Full-height Opteron server blades can take up to four dual core CPUs and Xeon server blades can take up to two quad cores.

The following HP ProLiant hardware models are supported for use in an HP XC hardware configuration:

- HP ProLiant BL460c (half-height)
  - Up to two quad core or dual core Intel® Xeon® processors
  - Two built-in network interface cards (NICs)
  - Two hot plug drives
  - Two mezzanine slots
- HP ProLiant BL465c (half-height)
  - Up to two single or dual core AMD® Opteron® processors
  - Two built-in network interface cards (NICs)
  - Two hot plug drives
  - Two mezzanine slots
- HP ProLiant BL480c (full-height)
  - Up to two quad core or dual core Xeon processors
  - Four built-in NICs
  - Four hot plug drives
  - Three mezzanine slots

- HP ProLiant BL685c (full-height)
  - Up to four dual core Opteron processors
  - Four built-in NICs
  - Two hot plug drives
  - Three mezzanine slots
- HP ProLiant BL860c (full-height)

1.4.2 Enclosures and Onboard Administrators

HP Server Blade c7000 Enclosure The HP Server Blade c7000 Enclosure is the enclosure model supported for use in an HP XC hardware configuration. An enclosure is a chassis that houses and connects blade hardware components. It can house a maximum of 16 half-height or 8 full-height server blades and contains a maximum of 6 power supplies and 10 fans.

The following are general guidelines for configuring enclosures:

- Up to four enclosures can be mounted in a 42U rack.
- If an enclosure is not fully populated with fans and power supplies, see the positioning guidelines in the HP Server Blade c7000 Enclosure documentation.
- Enclosures are cabled together using their uplink and downlink ports.
- The top uplink port in each rack is used as a service port to attach a laptop or other device for initial configuration or subsequent debugging .

The following enclosure setup guidelines are specific to HP XC:

- On every enclosure, an Ethernet interconnect module (either a switch or pass-thru module) is installed in bay 1 for the administration network.
- Hardware configurations that use Gigabit Ethernet as the interconnect require an additional Ethernet interconnect module (either a switch or pass-thru module) to be installed in bay 2 for the interconnect network.
- Systems that use InfiniBand as the interconnect require a double-wide InfiniBand interconnect switch module installed in double-wide bay 5 and 6.
- Some systems might need an additional Ethernet interconnect module to support server blades that require external connections. For more information about external connections, see "Cabling for the External Network" (page 19).

HP BladeSystem c-Class Onboard Administrator The Onboard Administrator is the management device for an enclosure, and at least one Onboard Administrator is installed in every enclosure. You can access the Onboard Administrator through a graphical Web-based user interface, a command-line interface, or the simple object access protocol (SOAP) to configure and monitor the enclosure. You can add a second Onboard Administrator to provide redundancy.

Insight Display The Insight Display is a small LCD panel on the front of an enclosure that provides instant access to important information about the enclosure such as the IP address and color-coded status. You can use the Insight Display panel to make some basic enclosure settings.

For more information about enclosures and their related components, see the *HP Server Blade c7000 Enclosure Setup and Installation Guide* .

1.4.3 iLO2 Console Management Device

Each HP ProLiant server blade has a built-in Integrated Lights Out (iLO2) device that provides full remote power control and serial console access. You can access the iLO2 device through the

Onboard Administrator. On server blades, iLO2 advanced features are enabled by default and include the following:

- Full remote graphics console access including full keyboard, video, mouse (KVM) access through a Web browser
- Support for remote virtual media which enables you to mount a local CD or diskette and serve it to the server blade over the network

1.4.4 Management Processor Console Management Device

Each HP Integrity server blade has a built-in management processor (MP) device that provides full remote power control and serial console access. You can access the MP device by connecting a serial terminal or laptop serial port to the local IO cable that is connected to the server blade.

1.4.5 Mezzanine Cards

The mezzanine slots on each server blade provide additional I/O capability. Mezzanine cards are PCI-Express cards that attach inside the server blade through a special connector and have no physical I/O ports on them. Card types include Ethernet, fibre channel, or 10 Gigabit Ethernet.

1.4.6 Interconnect Modules

An interconnect module provides the physical I/O for the built-in NICs or the supplemental mezzanine cards on the server blades. An interconnect module can be either a switch or a pass-thru module.

A switch provides local switching and minimizes cabling. Switch models that are supported as interconnect modules include, but are not limited to:

- Nortel GbE2c Gigabit Ethernet switch
- Cisco Catalyst Gigabit Ethernet switch
- HP 4x DDR InfiniBand switch
- Brocade SAN switch

A pass-thru module provides direct connections to the individual ports on each node and does not provide any local switching.

Bays in the back of each enclosure correspond to specific interfaces on the server blades. Thus, all I/O devices that correspond to a specific interconnect bay must be the same type.

Interconnect Bay Port Mapping

Connections between the server blades and the interconnect bays are hard wired. Each of the 8 interconnect bays in the back of the enclosure has a connection to each of the 16 server bays in the front of the enclosure. The built-in NIC or mezzanine card into which the interconnect blade connects depends on which interconnect bay it is plugged into. Because full-height blades consume two server bays, they have twice as many connections to each of the interconnect bays.

See the *HP BladeSystem Onboard Administrator User Guide* for illustrations of interconnect bay port mapping connections on half- and full-height server blades.

2 Task Summary and Checklist

This chapter contains a summary of the steps required to configure HP server blades in an HP XC cluster.

2.1 Best Practice for System Configuration

In order to function properly as an HP XC system, each component must be configured according to HP XC guidelines. To make configuration settings on certain components, an active network is required. However, on an HP XC system, the internal administration network is not operational until the head node is installed and running. Therefore, HP recommends that you install and configure the head node first and then, use the live administration network to make the configuration settings for the rest of the hardware components in the system.

Thus, the high level sequence of events is this:

1. Physically set up the enclosures, populate the enclosures with nodes, and cable all hardware components together.
2. Prepare the head node and any non-blade server nodes.
3. Install the HP XC System Software on the head node.
4. Run the `cluster_prep` command on the head node.
5. Run the `discover` command to discover the network components.
6. Connect to each Onboard Administrator and make required settings.
7. Run the `discover` command to discover the enclosures.
8. Run the `discover` command to discover the nodes.
9. Access all Onboard Administrators, console management devices (iLO2 or MP) and make required BIOS settings.

2.2 Installation and Configuration Checklist

Table 2-1 provides a checklist of tasks.

IMPORTANT: Hardware preparation is the key element in the successful installation and configuration of the system. If you do not prepare the hardware as described in this document, do not expect the cluster configuration process to be successful.

Table 2-1 Installation and Configuration Checklist

Task Category	Task Description	Where Documented
Cabling	Cable the switches and enclosures to configure the XC networks	Chapter 3
Software Installation	Gather the information you need for the installation	Section 4.1
	Install the HP XC system software on the head node	Section 4.2
Discovery	Gather information for the cluster preparation and discover process.	Section 5.1
	Optionally change the default IP address base of the HP XC networks	Section 5.2
	Run the `cluster_prep --enclosurebased` command on the head node	Section 5.3

Table 2-1 Installation and Configuration Checklist *(continued)*

Task Category	Task Description	Where Documented
	Run the `discover --enclosurebased --network` command on the head node to discover the switches	Section 5.4
	Set the Onboard Administrator password, which must match the passwords on the ProCurve switches and console management devices	Section 5.5
	Run the `discover --enclosurebased --enclosures` and `discover --enclosurebased --nodes` commands to discover the remainder of the hardware components	Section 5.6
BIOS Settings on Nodes	Make BIOS settings on non-blade servers according to regular procedures	Section 6.1 and the *HP XC Hardware Preparation Guide*
	Make BIOS settings on server blades	Section 6.2 and Section 6.3
System Configuration	Install software patches that might be available for the release of HP XC System Software you are installing	Section 7.1
	Perform various configuration tasks to set up the system environment	Section 7.2
	Run the `cluster_config` utility on the head node to configure the system and create the golden image	
	Run the `startsys` command to start all nodes in the system and propagate the golden image to all nodes	Section 7.2.2
	Run system verification tasks, including the operation verification program (OVP), to verify that the system is operating correctly	Section 7.3

3 Cabling

The following topics are addressed in this chapter:

- "Network Overview " (page 15)
- "Cabling for the Administration Network" (page 15)
- "Cabling for the Console Network" (page 16)
- "Cabling for the Interconnect Network" (page 17)
- "Cabling for the External Network" (page 19)

3.1 Network Overview

An HP XC system consists of several networks: administration, console, interconnect, and external (public). In order for these networks to function, you must connect the enclosures, server blades, and switches according to the guidelines provided in this chapter.

The *HP XC Hardware Preparation Guide* guide describes specific instructions about which ports on each ProCurve switch are used for specific node connections on non-blade server nodes.

A hardware configuration with server blades does not have these specific cabling requirements; specific switch port assignments are not required. However, HP recommends a logical ordering of the cables on the switches to facilitate serviceability. Enclosures are discovered in port order, so HP recommends that you cable them in the order you want them to be numbered. Also, HP recommends that you cable the enclosures in lower ports and cable the external nodes in the ports above them.

Appendix A (page 43) provides several network cabling illustrations based on the interconnect type and server blade height to use as a reference.

3.2 Cabling for the Administration Network

The HP XC administration network is a private network within an HP XC system that is used primarily for administrative operations. The administration network is created and connected through ProCurve model 2800 series switches. One switch is designated as the root administration switch and that switch can be connected to multiple branch administration switches, if required.

NIC1 on each server blade is dedicated as the connection to the administration network. NIC1 of all server blades connects to interconnect bay 1 on the enclosure.

The entire administration network is formed by connecting the device (either a switch or a pass-thru module) in interconnect bay 1 of each enclosure to one of the ProCurve administration network switches.

Non-blade server nodes must also be connected to the administration network. See the *HP XC Hardware Preparation Guide* to determine which port on the node is used for the administration network; the port you use depends on your particular hardware model.

Figure 3-1 illustrates the connections that form the administration network.

Figure 3-1 Administration Network Connections

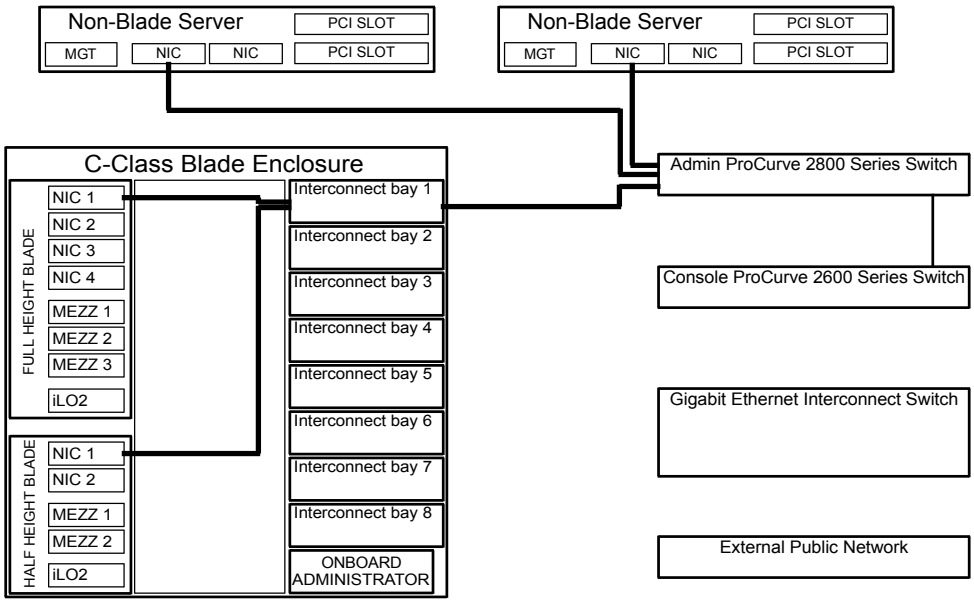

3.3 Cabling for the Console Network

The console network is part of the private administration network within an HP XC system, and it is used primarily for managing and monitoring the node consoles.

On a small cluster, the console management devices can share a single top-level ProCurve 2800 root administration switch. On larger hardware configurations that require more ports, the console network is formed with separate ProCurve model 2600 series switches.

You arrange these switches in a hierarchy similar to the administration network. One switch is designated as the root console switch and that switch can be connected to multiple branch console switches. The top-level root console switch is then connected to the root administration switch.

HP server blades use iLO2 as the console management device. Each iLO2 in an enclosure connects to the Onboard Administrator. To form the console network, connect the Onboard Administrator of each enclosure to one of the ProCurve console switches.

Non-blade server nodes must also be connected to the console network. See the *HP XC Hardware Preparation Guide* to determine which port on the node is used for the console network; the port you use depends on your particular hardware model.

Figure 3-2 illustrates the connections that form the console network.

Figure 3-2 Console Network Connections

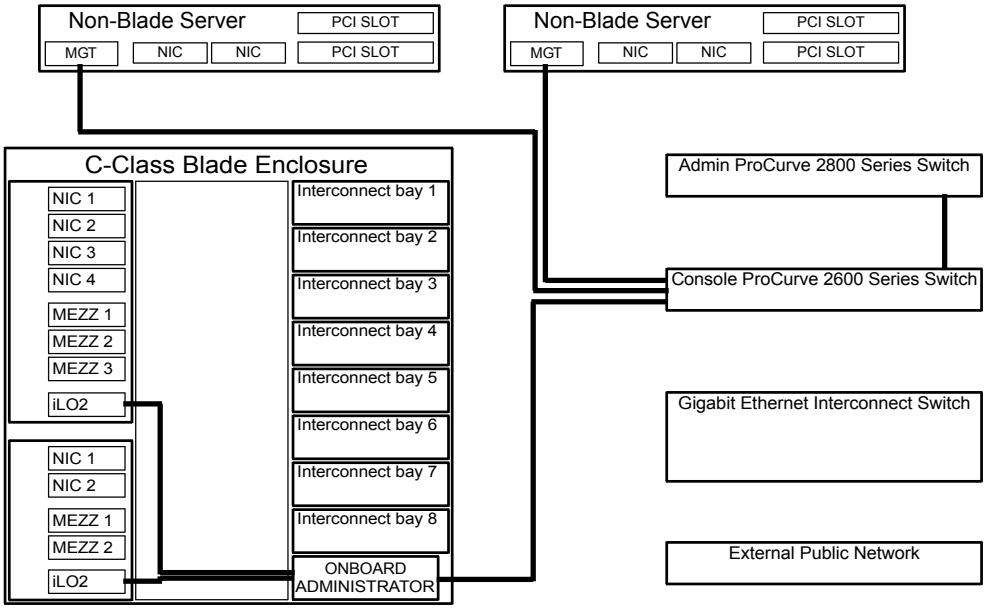

3.4 Cabling for the Interconnect Network

The interconnect network is a private network within an HP XC system. Typically, every node in an HP XC system is connected to the interconnect. The interconnect network is dedicated to communication between processors and access to data in storage areas. It provides a high-speed communications path used primarily for user file service and for communications within applications that are distributed among nodes in the cluster.

Gigabit Ethernet and InfiniBand are supported as the interconnect types for HP XC hardware configurations with server blades and enclosures. The procedure to configure the interconnect network depends upon the type of interconnect in use.

- "Configuring a Gigabit Ethernet Interconnect" (page 17)
- "Configuring an InfiniBand Interconnect" (page 18)
- "Configuring the Interconnect Network Over the Administration Network" (page 19)

3.4.1 Configuring a Gigabit Ethernet Interconnect

A Gigabit Ethernet interconnect requires one or more external Ethernet switches to act as the interconnect between the enclosures that make up the HP XC system.

On systems using a Gigabit Ethernet interconnect, one NIC on each server blade is dedicated as the connection to the interconnect network. On a server blade, NIC2 is used for this purpose. NIC2 of all server blades connects to interconnect bay 2 on the enclosure.

The entire interconnect network is formed by connecting the device (either a switch or a pass-thru module) in interconnect bay 2 of each enclosure to one of the Gigabit Ethernet interconnect switches.

If the device is a switch, the Gigabit uplink to the higher level ProCurve switch can be a single wire or a trunked connection of 2, 4, or 8 wires. If the device is a pass-thru module, there must be one uplink connection for each server blade in the enclosure.

Non-blade server nodes must also be connected to the interconnect network. See the *HP XC Hardware Preparation Guide* to determine which port on the node is used for the interconnect network; the port you use depends on your particular hardware model.

Figure 3-3 illustrates the connections for a Gigabit Ethernet interconnect.

Figure 3-3 Gigabit Ethernet Interconnect Connections

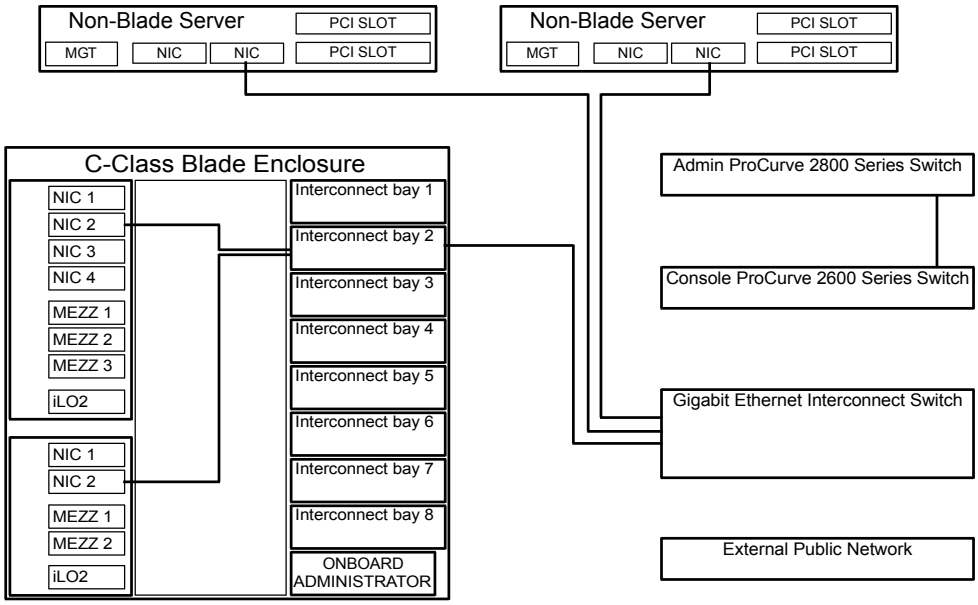

3.4.2 Configuring an InfiniBand Interconnect

An InfiniBand interconnect requires one or more external InfiniBand switches with at least one managed switch to manage the fabric.

Systems using an InfiniBand interconnect require you to install an InfiniBand mezzanine card into mezzanine bay 2 of each server blade to provide a connection to the InfiniBand interconnect network. The InfiniBand card in mezzanine bay 2 connects to the double-wide InfiniBand switch in interconnect bays 5 and 6 on the enclosure.

The entire interconnect network is formed by connecting the InfiniBand switches in interconnect bays 5 and 6 of each enclosure to one of the InfiniBand interconnect switches.

Non-blade server nodes also require InfiniBand cards and must also be connected to the interconnect network.

Figure 3-4 illustrates the connections for an InfiniBand interconnect.

Figure 3-4 InfiniBand Interconnect Connections

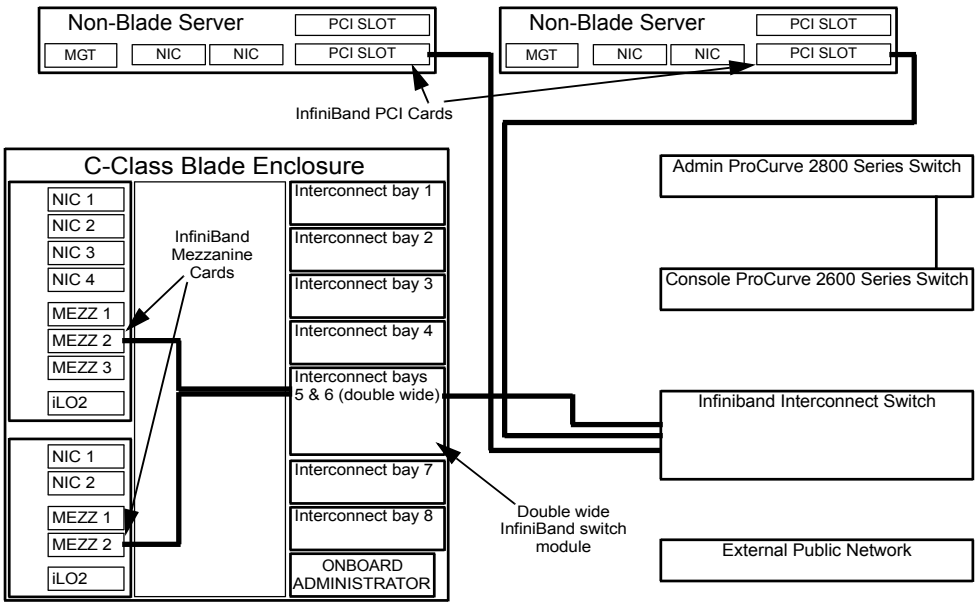

3.4.3 Configuring the Interconnect Network Over the Administration Network

In cases where an additional Gigabit Ethernet port or switch may not be available, the HP XC System Software enables you to configure the interconnect on the administration network. When the interconnect is configured on the administration network, only a single LAN is used.

To configure the interconnect on the administration network, you include the `--ic=AdminNet` option on the `discover` command line, which is documented in "Task 4: Discover Switches" (page 29).

Be aware that configuring the interconnect on the administration network may negatively impact system performance.

3.5 Cabling for the External Network

Depending upon the roles you assign to nodes during the cluster configuration process, some nodes might require connections to an external public network. Making these connections requires one or more Ethernet ports in addition to the ports already in use. The ports you use depend upon the hardware configuration and the number of available ports.

On non-blade server nodes, the appropriate port assignments for the external network are shown in the *HP XC Hardware Preparation Guide*.

On a server blade, the number of available Ethernet ports is influenced by the type of interconnect and the server blade height:

- Nodes in clusters that use an InfiniBand interconnect have only one NIC in use for the administration network.
- Nodes in clusters that use a Gigabit Ethernet interconnect have two NICs in use; one for the administration network, and one for the interconnect network.
- Half-height server blade models have two built-in NICs.
- Full-height server blade models have four built-in NICs.

You can use the built-in NICs on a server blade if any are available. If the node requires more ports, you must add an Ethernet card to mezzanine bay 1 on the server blade. If you add an Ethernet card to mezzanine bay 1, you must also add an Ethernet interconnect module (either a switch or pass-thru module) to interconnect bay 3 or 4 of the enclosure.

On full-height server blades, you can avoid having to purchase an additional mezzanine card and interconnect module by creating virtual local area networks (VLANs). On a full-height server blade, NICs 1 and 3 are both connected to interconnect bay 1, and NICs 2 and 4 are both connected to interconnect bay 2. If you are using one of these NICs for the connection to the external network, you might have to create a VLAN on the switch in that bay to separate the external network from other network traffic.

For information about configuring VLANs, see "Creating VLANs" (page 22).

The ports and interconnect bays used for external network connections vary depending on the hardware configuration, the ports that are already being used for the other networks, and the server blade height. For more information about how to configure the external network in these various configurations, see the illustrations in the following sections:

- "Configuring the External Network: Option 1" (page 19)
- "Configuring the External Network: Option 2" (page 20)
- "Configuring the External Network: Option 3 - Non Gigabit Ethernet Interconnect Clusters" (page 21)

3.5.1 Configuring the External Network: Option 1

Figure 3-5 (page 20) assumes that NIC1 and NIC2 are already in use for the administration and interconnect networks. This situation requires a third NIC for the external network. Half-height

server blades do not have three NICs, and therefore, half-height server blades are not included in this example

Because NIC1 and NIC3 on a full-height server blade are connected to interconnect bay 1, you must use VLANs on the switch in that bay to separate the external network from the administration network.

Also, in this example, PCI Ethernet cards are used in the non-blade server nodes. If the hardware configuration contains non-blade server nodes, see the *HP XC Hardware Preparation Guide* for information on which port to use for the external network.

Figure 3-5 External Network Connections: Full-Height Server Blades and NIC1 and NIC2 in Use

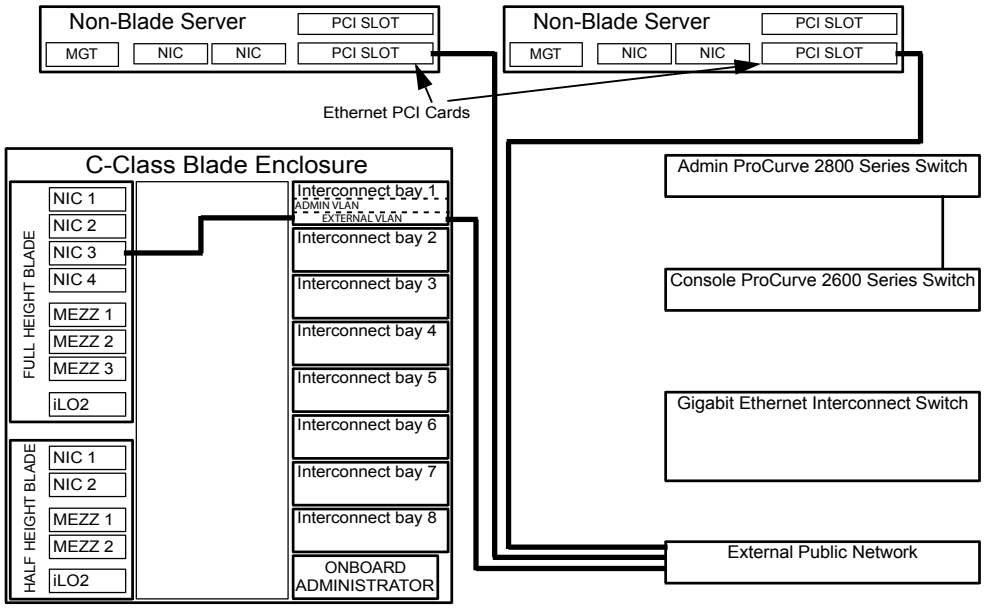

3.5.2 Configuring the External Network: Option 2

Figure 3-6 (page 21) assumes that NIC1 and NIC2 are already in use for the administration and interconnect networks. This situation requires a third NIC for the external network, but unlike Figure 3-5 (page 20), this hardware configuration includes half-height server blades. Therefore, to make another Ethernet NIC available, you must add an Ethernet card to mezzanine bay 1 on each server blade that requires an external connection. You must also install an Ethernet interconnect module in interconnect bay 3 for these cards.

In addition, PCI Ethernet Cards are used in the non-blade server nodes. If the hardware configuration contains non-blade server nodes, see the *HP XC Hardware Preparation Guide* for information on which port to use for the external network.

Figure 3-6 External Network Connections: Half-Height Server Blades and NIC1 and NIC2 in Use

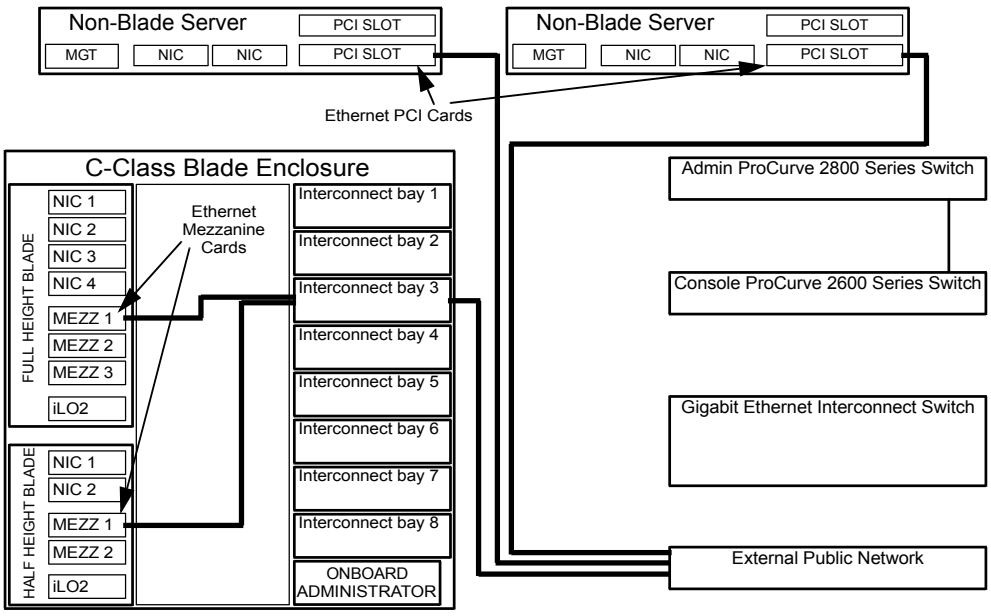

3.5.3 Configuring the External Network: Option 3 - Non Gigabit Ethernet Interconnect Clusters

The administration network requires only one network interface, NIC1, on clusters that do not use Gigabit Ethernet as the interconnect (that is, they use InfiniBand or the interconnect on the administration network).

On these non Gigabit Ethernet interconnect clusters, you have two configuration methods to configure an external network connection, and the option you choose depends on whether the collection of nodes requiring external connections includes half-height server blades.

- If only full height server blades require external connections, you can use NIC3 for the external network. This is similar to the way the external connection is configured in Figure 3-5 (page 20), and it saves the cost of an additional interconnect device in bay 2.
- If half-height server blades require external connections, you cannot use NIC3 because half height server blades do not have a third NIC. In this case, you must use NIC2 as the external connection as shown in Figure 3-5 (page 20). This configuration requires an Ethernet interconnect module to be present in bay 2.

Figure 3-5 (page 20) also shows the use of built-in NICs in the non-blade server nodes for the external connection, but this varies by hardware model.

If the hardware configuration contains non-blade server nodes, see the *HP XC Hardware Preparation Guide* for information about which port to use for the external network.

Figure 3-7 External Network Connections: Half and Full-Height Server Blades and NIC1 in Use

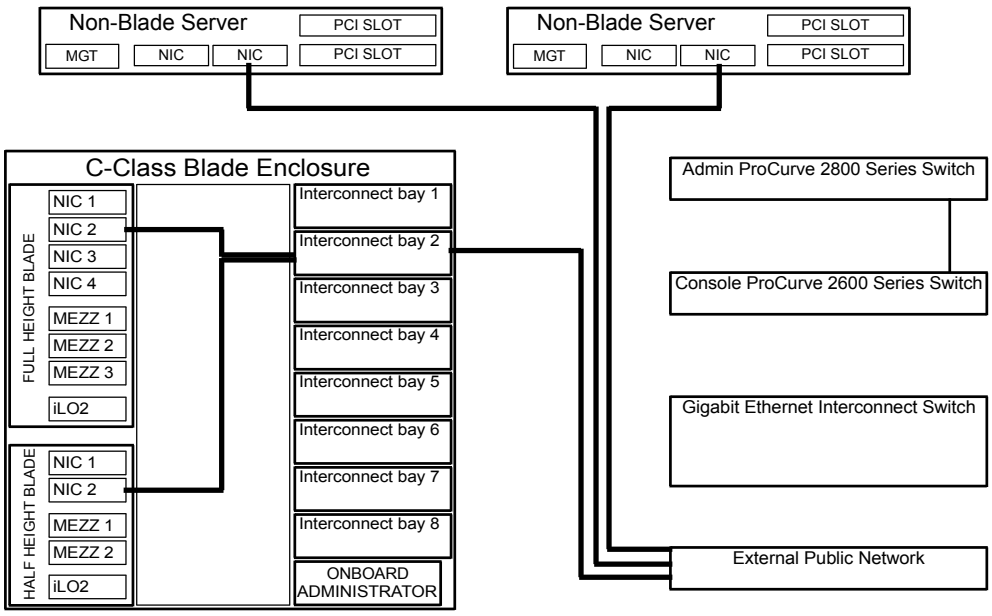

3.5.4 Creating VLANs

Use the following procedure on GbE2c (Nortel) switches if you need to configure a VLAN to separate the external network from other network traffic.

1. See the illustrations of interconnect bay port mapping connections in the *HP BladeSystem Onboard Administrator User Guide* to determine which ports on the switch to connect to each of the two virtual networks. Remember to include at least one of the externally accessible ports in each VLAN.
2. Connect a serial device to the serial console port of the GbE2c switch.
3. Press the **Enter** key.
4. When you are prompted for a password, enter **`admin`**, which is the default password.
5. Enter the following commands to access the VLAN configuration:
   a. **`cfg`**
   b. **`l2`** (the letter **`l`** as in **`l`**ayer, not the number one)
   c. **`vlan 2`** (be sure to enter a space between **`vlan`** and the vlan number)
6. Specify a name for the VLAN; choose any name you want.

   ```
   # name your_name
   ```
7. Enable the VLAN:

   ```
   # ena
   ```
8. Add each port to the VLAN one at a time. If you see a message that the port is in another VLAN, answer **yes** to move it. This example adds ports 1, 3, and 21 to the VLAN.

   ```
   # add 1
   # add 3
   # add 21
   ```

   If you need more information about creating VLANs, see the GbE2c documentation.
9. When you have completed adding ports, enter **`apply`** to activate your changes and enter **`save`** to save them.

4 Installing HP XC System Software On the Head Node

The information in this chapter parallels the information in Chapter 2 in the *HP XC System Software Installation Guide*. At some points, you might be instructed to refer to that document.

Table 4-1 lists where to find head node installation instructions. The instructions differ if the head node is an HP server blade.

Table 4-1 Head Node Installation Instructions

Is the Head Node a Server Blade?	Installation Procedure
No	1. Follow the instructions in the *HP XC Hardware Preparation Guide* to make the appropriate BIOS settings on the head node. 2. Then, follow the instructions in Chapter 2 of the *HP XC System Software Installation Guide* to install the HP XC System Software on the head node.
Yes	Follow the procedures in this chapter: 1. "Task 1: Refer to the Installation Guide" (page 23) 2. "Task 2: Install HP XC System Software on the Head Node" (page 23)

4.1 Task 1: Refer to the Installation Guide

Open the *HP XC System Software Installation Guide* , start reading at Chapter 1, and continue reading and performing tasks until you reach Section 2.3 in Chapter 2. Stop there and return to this HowTo.

4.2 Task 2: Install HP XC System Software on the Head Node

A server blade does not have a local DVD drive, a local VGA connection, a local keyboard, or mouse ports. Thus, you must install the HP XC System Software using the virtual console and virtual media features of the iLO2 console management device. You cannot use the installation procedure in the *HP XC System Software Installation Guide*.

The virtual console interface enables you to connect to an iLO2 device over the network and display a full screen VGA graphics display. The virtual media function enables you to mount a CD or DVD drive on your local PC or laptop as if it were attached to the server blade.

Before you begin, you must have the following items available:

- A PC or laptop with a local DVD drive
- A recent version of Mozilla Firefox or Internet Explorer running on the PC or laptop

When you have these items in your possession, proceed to "Connect to the Onboard Administrator" (page 23).

4.2.1 Connect to the Onboard Administrator

Network access to the Onboard Administrator and iLO2 associated with the head node is required to use the virtual media features. The internal administration network is not operational until the head node is installed. Therefore, you must use one of the following methods to make the Onboard Administrator network accessible. If you are not able to connect to a public network (as described in the first method), use the second method.

Method 1

Put the Onboard Administrator on an active network:

1. To provide IP addresses for the Onboard Administrator and all iLO2 devices in the enclosure, disconnect the Onboard Administrator associated with the head node from the administration network, and plug the Onboard Administrator into an active network with a DHCP server.
2. Obtain the IP address of the Onboard Administrator from the Insight Display panel on the enclosure.
3. From your PC or laptop, use the browser to access the Onboard Administrator using the IP address you obtained in the previous step.
4. Log in to the Onboard Administrator with the default user name `Administrator` and the default password shown on the tag affixed to the Onboard Administrator.
5. Proceed to "Start the Installation (page 24)."

Method 2

Assign a temporary static IP address to the Onboard Administrator:

1. Use a network cable to connect the laptop or PC NIC directly into the Onboard Administrator or connect it to the administration network ProCurve switch.
2. Use the menus on the Insight Display panel to manually set a static IP address and subnet mask for the Onboard Administrator. You can use any valid IP address because there is no connection to a public network.

   All static addresses must be in the same network. For example, assume the network is 172.100.100.0 and the netmask is 255.255.255.0. In this case, the static IP addresses might be:

   - IP address of the installation PC: 172.100.100.20
   - IP address of the Onboard Administrator: 172.100.100.21
   - Starting IP address for enclosure bay IP addressing: 172.100.100.1 (this uses the addresses from 172.100.100.1 to 172.100.100.16)
3. On your laptop or PC, manually set a static IP address for the NIC in the same subnet as the IP address you just set for the Onboard Administrator.
4. From your PC or laptop, use the browser to access the Onboard Administrator using the static IP address you assigned to the Onboard Administrator.
5. Log in to the Onboard Administrator. Use the default user name `Administrator` and the default password shown on the tag affixed to the Onboard Administrator.
6. Assign IP addresses to all iLO2 devices in the enclosure:
   a. Click on the plus sign (+) to open the **Enclosure Settings** menu in the left frame.
   b. Select the **Enclosure Bay IP Addressing** menu item.
   c. Select the check box to **Enable Enclosure Bay IP Addressing**.
   d. Specify a beginning **IP Address** and **Subnet Mask** in the same subnet as the static address you set for the Onboard Administrator. The enclosure consumes 16 addresses starting at the address you specify, so make sure none of these addresses conflict with the two IP address you have already assigned.
   e. Click the **Apply** button to save your settings.
7. Proceed to "Start the Installation (page 24)."

4.2.2 Start the Installation

Follow this procedure to install the HP XC System Software on a server blade head node; this procedure assumes that you are logged in to the Onboard Administrator:

1. Insert the HP XC System Software Version 3.1 or Version 3.2 DVD into the DVD drive of the local laptop or PC.
2. Do the following from the Onboard Administrator management Web page:
   a. In the left frame, click the plus sign (+) to open the **Device Bays** menu.
   b. Click the plus sign (+) next to the node that represents the head node.
   c. Click on the link to the **iLO** to open the Integrated Lights-Out 2 Web utility.
3. Click the **Web Administration** link in the body of the window to open the iLO2 Web Administration page.
4. Click the **Remote Console** tab.
5. Open the remote console, map the local DVD drive to the server blade, and turn on power to the node. The procedure depends upon the type of browser you are using:

Internet Explorer

If you are using Internet Explorer as your browser, do the following:

a. Click the **Integrated Remote Console** link to open a virtual console window that provides access to the graphics console, virtual media, and power control in the same window.

b. Click the **Drives** pull-down menu.

c. Click **Mount** to mount the appropriate DVD drive

d. Click the **Power** button and then click **Momentary Press** to turn on power to the server and start booting from the DVD.

e. Proceed to step 6.

Mozilla Firefox

If you are using Firefox as your browser, do the following:

a. Click the **Remote Console** link to open the virtual console window.

b. In the iLO2 Web Administration window, click the **Virtual Devices** tab.

c. In the left frame, click the **Virtual Media** link.

d. Click the **Virtual Media Applet** link to open a new window that provides access to the virtual media functions.

e. Go to Virtual Media Applet window, and in the Virtual CD-ROM area of the window, click **Local Media Drive** and use the drop-down menu to select the drive where you inserted the HP XC DVD.

f. In the left frame of the iLO2 Web Administration window, click the **Virtual Power** link to access the power functions and then click **Momentary Press** to turn on power to the server and start booting from the DVD.

g. Proceed to step 6.

6. Look at the remote console window and watch as the XC DVD boots.
7. Start the Kickstart installation process when the `Boot:` prompt is displayed:

   ```
   Boot: linux ks=hd:scd0/ks.cfg
   ```

   Some hardware models require additional options to be included on the boot command line. Before booting the head node, look in Table 4-2 to determine if your head node requires a special option.

Table 4-2 Boot Command Line Options Based on Hardware Model

Hardware Model	Boot Command Line
HP ProLiant BL465c	`boot: linux ks=hd:scd0/ks.cfg pci=nommconf`
HP ProLiant BL685c	`boot: linux ks=hd:scd0/ks.cfg pci=nommconf`

8. See Table 2-5 in the *HP XC System Software Installation Guide*, which describes each installation prompt and provides information to help you with your answers.
9. When the HP XC software installation process is complete, log in as the root user.
10. Follow the instructions in Sections 2.4 and 2.5 in the *HP XC System Software Installation Guide* to install additional third-party software on the head node, if required.

5 Discovering the Hardware Components

This chapter describes the following tasks, which you must complete in the order shown:

- "Task 1: Prepare for the System Configuration" (page 27)
- "Task 2: Change the Default IP Address Base (Optional)" (page 28)
- "Task 3: Use the cluster_prep Command to Prepare the System" (page 28)
- "Task 4: Discover Switches" (page 29)
- "Task 5: Set the Onboard Administrator Password" (page 30)
- "Task 6: Discover Enclosures and Nodes" (page 30)

5.1 Task 1: Prepare for the System Configuration

Open the *HP XC System Software Installation Guide* to Chapter 3, and follow the instructions in Task 1, where you gather the information required for the system preparation and discovery phase.

Read the information in "Node Naming Differences" in this HowTo, which describes internal node numbering conventions for server blades that are not described in the Chapter 3 in the *HP XC System Software Installation Guide*.

5.1.1 Node Naming Differences

As described in Table 3-1 in the *HP XC System Software Installation Guide*, internal node naming differs when the hardware configuration contains enclosures and HP server blades.

When the hardware configuration does not contain enclosures and HP server blades, internal node names are assigned in a dense fashion in which there are no missing numbers in the node numbering scheme except for a possible missing number between the branch nodes and those nodes that are connected to the root administration switch.

In an enclosure-based system, the `discover` command uses a sparse node numbering scheme. This means that internal node names are assigned based on the enclosure in which the node is located and the bay the node is plugged into.

For example, if a node is plugged into bay 10 of enclosure 1, the node is numbered {*node_prefix*}10. In a configuration with two enclosures in which there might be 16 nodes in each enclosure, the node in bay 10 in enclosure 2 is numbered {*node_prefix*}27.

In this release, 16 is the maximum number of server blade nodes in a **real enclosure**. A real enclosure is defined as an enclosure with one of more Onboard Administrators. The maximum number of non-blade server nodes in a **virtual enclosure** is 38. A virtual enclosure is defined as a ProCurve switch that has at least one console port from a non-blade server node plugged into it.

5.1.2 Head Node Naming

When the hardware configuration does not contain (real) enclosures and HP server blades, the `cluster_prep` command determines and assigns the head node name by using the number that represents the maximum number of nodes allowed in the hardware configuration, which you supply. For example, if the maximum number of nodes allowed in your hardware configuration is 128, the head node is node {*node_prefix*}128.

When the hardware configuration contains server blades, the head node is named for its location in the system just like every other node, as described in "Node Naming Differences". The only exception is when the head node is a non-blade server node whose console port is not connected to the administration network ProCurve switch. In this case, the head node is named {*node_prefix*}0.

5.2 Task 2: Change the Default IP Address Base (Optional)

You may need to change the default IP address base of the HP XC private networks if the default values conflict with another IP address base at your site. The IP address base is defined in the `base_addrV2.ini` file.

Follow this optional procedure to change the default IP address base for the HP XC private networks:

1. Begin this procedure as the root user on the head node.
2. Use the text editor of your choice to edit the `/opt/hptc/config/base_addrV2.ini` file.

   The default file content is similar to this:

```
# This file contains the base addresses for the network components
# in XC cluster.
[global]
        nodeBase=172.31.0
        cpBase=172.31.16
        swBase=172.31.32
        netMask=255.255.192.0
        icBase=172.31.64
        icNetMask=255.255.192.0
        dyStart=172.31.48.1
        dyEnd=172.31.63.254
```

   The following describes the parameters in this file:

`nodeBase`	The base IP address of all nodes.
`cpBase`	The base IP address of the console branch.
`swBase`	The base IP address of the Ethernet switches. The IP address range for InfiniBand interconnect switch modules is based on this value, `172.31.33.*`.
`netMask`	The netmask that is used for all connections on the administration network.
`icBase`	The base IP address of the interconnect.
`icNetMask`	The base network mask of the interconnect.
`dyStart`	The start of the dynamic IP address range used during the discovery process.
`dyEnd`	The end of the dynamic IP address range used during the discovery process.

3. Follow these guidelines if you make changes to the `base_addrV2.ini` file:
   - Change only the first octet of a base address; 172 is the default.
   - Use the same first octet for all base IP addresses.
   - Use a private domain as defined by the RFC.
4. Save your changes and exit the text editor. Because this is a read-only file, the method by which you save the file depends upon the text editor you used.

   For example, if you are using the `vi` text editor, enter **`:w!`** to save changes to a read-only file.
5. Repeat steps 1 through 4 to edit the `base_addr.ini` file and make the same changes you made to the `base_addrV2.ini` file.

Proceed to "Task 3: Use the cluster_prep Command to Prepare the System".

5.3 Task 3: Use the cluster_prep Command to Prepare the System

Use the `cluster_prep` command to define the node naming prefix, to set the configuration and management database administrator's password, to configure the external Ethernet connection

on the head node, to start the MySQL service, and initialize the configuration and management database:

1. Begin this procedure as the root user on the head node.
2. Change to the following directory:

   ```
   # cd /opt/hptc/config/sbin
   ```

3. Start the system preparation process:

   ```
   # ./cluster_prep --enclosurebased
   ```

4. See Chapter 3, Task 3 in the *HP XC System Software Installation Guide*, which describes each `cluster_prep` prompt and provides information to help you with your answers.

   When the `cluster_prep` processing is complete, the head node is accessible through the external network.

5. To proceed with the discovery process, you **must** reset the temporary cabling and settings you made to the Onboard Administrator. Do one of the following, depending upon how you accessed the Onboard Administrator:
   - If you accessed the Onboard Administrator through a connection to a public network, return all cabling to its original configuration and press the reset button on the Onboard Administrator.
   - If you assigned a temporary static IP address to access the Onboard Administrator, do the following:
     a. Clear all enclosure bay IP addressing settings.
     b. Use the Insight Display panel on the enclosure to set the Onboard Administrator for DHCP.
     c. Return all cabling to its original configuration.
     d. Press the reset button on the Onboard Administrator.

6. Log in to the head node as root using the head node external IP address:

   ```
   # ssh IP_address
   ```

Proceed to "Task 4: Discover Switches."

5.4 Task 4: Discover Switches

The HP XC System Software has been designed to communicate directly with the Onboard Administrator on each enclosure. Thus, the `discover` command obtains the required configuration information directly from the Onboard Administrator to automatically discover all components in the hardware configuration.

Follow this procedure to discover the switches and activate the administration network. Before you begin, you must have already returned the Onboard Administrator to its original configuration and you must be logged in to the head node through the external network connection.

1. Change to the following directory:

   ```
   # cd /opt/hptc/config/sbin
   ```

2. Discover the switches:

   ```
   # ./discover --enclosurebased --network [--ic=AdminNet] --verbose
   ```

NOTE: The `[--ic=AdminNet]` option is required only if you are configuring the interconnect on the administration network.

Command output looks similar to the following:

```
Discovery - XC Cluster version HP XC V3.1 RC2 20061022
```

```
Enter the MAC address of the switch the head node plugs into
in the format xx:xx:xx:xx:xx:xx : your_MAC_address
Please enter the Procurve switch Administrator password : your_password
Please re-enter password: your_password

Discovering Switches...
Restarting dhcpd
Waiting for IP addresses to stabilize ..................... done

discoverSwitches: IP: 172.31.32.1 name: blc1nxcsw000000000000-0
model: 2824 MaxPorts: 24
Checking switch 172.31.32.1 for neighboring switches ... done

discoverSwitches: IP: 172.31.63.252 name: blc1nxcsw001708ad3740-24
model: 2626 MaxPorts: 26
Checking switch 172.31.63.252 for neighboring switches ... done

found switch blc1nxcsw000000000000-0 in both database and switch Hash,
ignoring
Restarting dhcpd waiting for network component at 172.31.32.2 to
become available.
........
switchName blc1nxcsw001708ad3740-24 switchIP 172.31.32.2 type 2626
switchName blc1nxcsw000000000000-0 switchIP 172.31.32.1 type 2824
```

Proceed to "Task 5: Set the Onboard Administrator Password."

5.5 Task 5: Set the Onboard Administrator Password

You must define and set the user name and password for the Onboard Administrator on every enclosure in the hardware configuration.

The Onboard Administrator user name and password must match the user name and password you plan to use for the iLO2 console management devices. The default user name is `Administrator`, and HP recommends that you delete the predefined `Administrator` user for security purposes.

If you are using the default user name `Administrator`, set the password to be the same as the iLO2. If you create a new user name and password for the iLO2 devices, you must make the same settings on all Onboard Administrators.

Follow this procedure to configure a common password for each active Onboard Administrator:

1. Use a network cable to plug in your PC or laptop to the administration network ProCurve switch.
2. Make sure the laptop or PC is set for a DHCP network.
3. Gather the following information:
   a. Look at the Insight Display panel on each enclosure, and record the IP address of the Onboard Administrator.
   b. Look at the tag affixed to each enclosure, and record the default Onboard Administrator password shown on the tag.
4. On your PC or laptop, use the information gathered in the previous step to browse to the Onboard Administrator for every enclosure, and set a common user name and password for each one. This password must match the administrator password you will later set on the ProCurve switches. Do not use any special characters as part of the password.

Proceed to "Task 6: Discover Enclosures and Nodes."

5.6 Task 6: Discover Enclosures and Nodes

Follow this procedure to discover all enclosures and nodes in the hardware configuration. This discovery process assigns IP addresses to all hardware components:

1. Change to the following directory:

   ```
   # cd /opt/hptc/config/sbin
   ```

2. Discover all enclosures:

   ```
   # ./discover --enclosurebased --enclosures

   Discovery - XC Cluster version HP XC V3.1 RC2 20061022

   Enter the common user name for all console port management devices:
    your_username
   Please enter the password for Administrator: your_password
   Please re-enter password: your_password
   Discovering blade enclosures ...
   Checking switch 172.31.32.2 for active ports ...done
   Getting MAC addresses from switch 172.31.32.2 ... done
   Checking switch 172.31.32.1 for active ports ...done
   Getting MAC addresses from switch 172.31.32.1 ... done
   Enclosure blc1n-enc09USE6391TF5 found
   Discovering virtual enclosures ...
   Checking switch 172.31.32.2 for active ports ... done
   Getting MAC Addresses from switch 172.31.32.2 ... done
   Checking switch 172.31.32.1 for active ports ... done
   Getting MAC Addresses from switch 172.31.32.1 ...done
   ```

3. Use one of the following options to discover all nodes. The command line depends on the version of HP XC System Software you are installing:

   - HP XC System Software Version 3.1:

     ```
     # ./discover --enclosurebased --nodes --verbose
     ```

   - HP XC System Software Version 3.2:

     ```
     # ./discover --enclosurebased --nodesonly --verbose
     ```

   Command output is similar to the following:

   ```
   Discovery - XC Cluster version HP XC Vn.n <timestamp>

   Enter the common user name for all console port management devices:
   your_username
   Please enter the password for Administrator: your_password
   Please re-enter password: your_password
   Discovering Blades ...
   Blade found in enclosure blc1n-enc09USE6391TF5 bay 1, name is blc1n1
   Blade found in enclosure blc1n-enc09USE6391TF5 bay 2, name is blc1n2
   Blade found in enclosure blc1n-enc09USE6391TF5 bay 3, name is blc1n3
   Blade found in enclosure blc1n-enc09USE6391TF5 bay 4, name is blc1n4
   Blade found in enclosure blc1n-enc09USE6391TF5 bay 5, name is blc1n5
   Discovering Non-Blade systems ...
   Setting system name to blc1n0
   uploading database
   Restarting dhcpd
   Opening /etc/hosts
   Opening /etc/hosts.new.XC
   Opening /etc/powerd.conf
   Building /etc/powerd.conf ... done
   Attempting to start hpls power daemon ... done
   Waiting for power daemon ... done

   Checking if all console ports are reachable ...
   number of cps to check, 5
   checking 172.31.16.5
   checking 172.31.16.1
   checking 172.31.16.4
   checking 172.31.16.3
   checking 172.31.16.2  .done
   ```

```
Starting CMF for discover...
Stopping cmfd:                                             [FAILED] 1
Starting cmfd:                                             [  OK  ]
Waiting for CMF to establish console connections .......... done

uploading database
Restarting dhcpd
Opening /etc/hosts
Opening /etc/hosts.new.XC
Opening /etc/powerd.conf
Building /etc/powerd.conf ... done
Attempting to start hpls power daemon ... done
Waiting for power daemon ... done

uploading database
```

1 The `discover` command turns off the console management facility (CMF) daemon. If the CMF daemon is not running, a "FAILED" message is displayed. This message is expected, and you can ignore it.

6 Making Node-Specific Settings

Running the discover process has assigned IP addresses to all hardware components. You can now use the active administration network to access the various hardware components to make node-specific BIOS settings that are required for HP XC.

IMPORTANT: Making the required settings on each node is a key element in the successful installation and configuration of the system. If you do not prepare the nodes as described in this chapter, do not expect the cluster configuration process to be successful.

This chapter describes the following tasks. You must perform these tasks on each node in your hardware configuration:

- "Making Settings on Non-Blade Servers" (page 33)
- "Making Settings on HP ProLiant Server Blades" (page 33)
- "Making Settings on HP Integrity Server Blades" (page 36)

6.1 Making Settings on Non-Blade Servers

Prepare non-blade server nodes according to the instructions in the *HP XC Hardware Preparation Guide*.

Network access to the iLO is required to enable `telnet` access and to turn off the automatic power-on mode. Network access is not possible until the iLO devices obtain their IP addresses from DHCP. However, when you follow the instructions in this HowTo, you have already run the `discover` command and all IP addresses have been assigned. Thus, you can ignore the instructions in the *HP XC Hardware Preparation Guide* that instruct you to go the *HP XC System Software Installation Guide* to obtain IP addresses.

6.2 Making Settings on HP ProLiant Server Blades

You must perform the following tasks on each server blade in the hardware configuration:

- Set the boot order
- Create an iLO2 user name and password
- Set the power regulator
- Configure smart array devices

You use the Onboard Administrator, the iLO2 web interface, and virtual media to make the appropriate settings on HP ProLiant Server Blades.

Setup Procedure

Perform the following procedure or each HP server blade in the hardware configuration:

1. On the head node, look in the `hosts` file to obtain the IP addresses of the Onboard Administrators:

   ```
   # grep OA1 /etc/hosts
   ```

2. Record the IP addresses for use later in this procedure.
3. Use a network cable to connect your PC or laptop to the ProCurve administration switch.
4. From a PC or laptop, use the browser to access an Onboard Administrator. In the **Address** field of the browser, enter one of the IP addresses you obtained in step 1, similar to the following:

   Address `https://172.31.32.2`

5. When the login window is displayed, log in to the Onboard Administrator with the user name and password you set previously.

6. In the left frame, click the plus sign (+) next to **Device Bays** to display the list of nodes contained in the enclosure.
7. Click the link to the first hardware model in the list. Wait a few seconds until the frame to the right is populated with node-specific information.
8. Click the **Boot Options** tab.
   a. Select a boot device, and use the up and down arrows on the screen to position the device so that it matches the boot order listed in Table 6-1.

Table 6-1 Boot Order for HP ProLiant Server Blades

Head Node	All Other Nodes
Set the following boot order on the head node: 1. USB 2. Floppy 3. CD 4. Hard Disk 5. PXE NIC1	Set the following boot order on all nodes except the head node: 1. USB 2. Floppy 3. CD 4. PXE NIC 1 5. Hard Disk

   b. Click the **Apply** button.
9. In the left frame, do the following to create a new iLO2 user name and password on this node:
   a. Under the hardware model, click **iLO**.
   b. In the body of the main window, click the **Web Administration** link to open the Integrated Lights-Out 2 utility in a new window. You might have to turn off popup blocking for this window to open.
   c. In the new window, click the **Administration** tab.
   d. In the left frame, click the **User Administration** link.
   e. Click the **New** button, and create a new iLO2 user name and password, which must match the user name and password you set on the Onboard Administrator. Do not use any special characters as part of the password.

      You use this user name and password whenever you need to access the console port with the `telnet cp-`*`nodename`* command.

   f. HP recommends that you delete the predefined `Administrator` user for security purposes.

      The Onboard Administrator automatically creates user accounts for itself (prefixed with the letters `OA`) to provide single sign-on capabilities. Do not remove these accounts.
10. Enable **telnet** access:
    a. In the left frame, click **Access**.
    b. Click the control to enable **Telnet Access**.
    c. Click the **Apply** button to save the settings.
11. Click the **Virtual Devices** tab and make the following settings:
    a. For every node except the head node, select **No** to **Automatically Power On Server** because you do not want to automatically turn on power to the node.
    b. Click the **Submit** button.
    c. In the left frame, click on the **Power Regulator** link.
    d. Do one of the following:
       - On Xeon-based server blades, select **Enable HP Static High Performance Mode**.
       - On Opteron-based server blades, select **Enable OS Control Mode**.

e. Click the **Apply** button to save the settings.

12. Configure disks into the smart array from the remote graphics console.

    All server blades have smart array cards, you must add the disk or disks to the smart array before attempting to image the node.

    To set up the smart array device, click the **Remote Console** tab on the virtual console page of the iLO2 Web Administration Utility, and then do one of the following depending on the browser type.

    #### Internet Explorer

    If you are using Internet Explorer as your browser, do the following:

    a. Click the **Integrated Remote Console** link to open a remote console window which provides access to the graphics console virtual media and power functions.
    b. In the remote console window, click the **Power** button.
    c. Click the **Momentary Press** button.
    d. Wait a few seconds for the power up phase to begin. Click the MB1 mouse button in the remote console window to put the pointer focus in this window so that your keyboard strokes are recognized.
    e. Proceed to step 13.

    #### Mozilla Firefox

    If you are using Mozilla Firefox as your browser, do the following:

    a. Click the **Remote Console** link to open a virtual console window.
    b. Go back to the iLO2 utility Web page and click the **Virtual Devices** tab.
    c. Click the **Momentary Press** button.
    d. Go back to the remote console window. Wait a few seconds for the power up phase to begin. Click the MB1 mouse button in this window to put the pointer focus in the remote console window so that your keyboard strokes are recognized in this window.
    e. Proceed to step 13.

13. Watch the screen carefully during the power-on self-test phase, and press the **F8** key when you are prompted to configure the disks into the smart array. Select **View Logical Drives** to determine if a logical drives exists. If a logical drive is not present, create one.

    If you create a logical drive, exit the SmartArray utility and power off the node. Do not let it try to boot up.

    Specific smart array configuration instructions are outside the scope of this HowTo. See the documentation that came with your model of HP ProLiant server for more information.

14. Perform this step for HP ProLiant BL685c nodes; proceed to the next step for all other hardware models.

    On an HP ProLiant BL685c node, watch the screen carefully during the power-on self-test, and press the **F9** key to access the ROM-Based Setup Utility (RBSU) to enable HPET as shown Table 6-2.

    Table 6-2 Additional BIOS Setting for HP ProLiant BL685c Nodes

Menu Name	Option Name	Set To This Value
Advanced	Linux x86_64 HPET	`Enabled`

    Press the **F10** key to exit the RBSU. The server automatically restarts.

15. Use the virtual power functions to turn off power to the server blade.

16. Close the iLO2 utility Web page.
17. Repeat this procedure from every active Onboard Administrator and make the same settings for each server blade in each enclosure.

6.3 Making Settings on HP Integrity Server Blades

Use the management processor (MP) on each HP Integrity server blade to make the following settings:

- Clear all event logs
- Enable IPMI over LAN
- Create an MP login ID and password that matches all other devices
- Add a boot entry for the string `DVD boot` on the head node, and add a boot entry for the string `Netboot` on all other nodes.
- Move the `DVD boot` and `Netboot` boot entries to top of the boot order
- Set the primary console
- Set the console baud rate to 115200

Procedure

Perform the following procedure on each HP Integrity server blade. All nodes must be seated in an enclosure, plugged in, and power turned off.

1. Connect the local IO cable (also called the SUV cable) to the server blade. The local cable is shipped with the enclosure and connects to the server blade at one end and is divided into VGA, USB, and Serial ports at the other end.
2. Connect a serial terminal or laptop serial port to the serial port of the local IO cable.
3. Use a terminal emulator, such as HyperTerminal, to open a terminal window.
4. Press the **Enter** key to access the MP. If there is no response, press the MP reset pin on the back of the MP and try again.
5. Log in to the MP using the default administrator name and password that shown on the screen. The **MP Main Menu** is displayed.

```
MP MAIN MENU:
CO: Console
VFP: Virtual Front Panel
CM: Command Menu
SMCLP: Server Management Command Line Protocol
CL: Console Log
SL: Show Event Logs
HE: Main Help Menu
X: Exit Connection
```

6. Enter **`SL`** to show event logs. Then, enter **`C`** to clear all log files and **`y`** to confirm your action.
7. Enter **`CM`** to display the **Command Menu**.
8. Do the following to ensure that the IPMI over LAN option is set. This setting is required for Nagios monitoring.
   a. Enter **`SA`** to display the **Set Access** configuration menu.
   b. Verify that the IPMI over LAN option is enabled.
   c. Enable the IPMI over LAN option if it is disabled.
      1. Enter the letter **`i`** to access the IPMI
      2. Enter the letter **`e`** to enable the IPMI over LAN
      3. Enter the letter **`y`** to confirm your action.
   d. Return to the **Command Menu**.

9. Enter **UC** (user configuration) and use the menu options to remove the default administrator and operator accounts. Then, for security purposes, create your own unique user login ID and password. Assign all rights (privileges) to this new user.

   The user login ID must have a minimum of 6 characters, and the password must have exactly 8 characters. You must set the same user login ID and password on every node and all MPs, iLOs, and OAs must use the same user name and password. Do not use any special characters as part of the password.

10. Turn on power to the node:

    ```
    MP:CM> pc -on -nc
    ```

11. Press **Ctrl+b** to return to the **MP Main Menu**.
12. Enter **CO** to connect to the console. It takes a few minutes for the live console to display.
13. Add a boot entry and set the OS boot order. Your actions for the head node differ from all other nodes.

Table 6-3 Adding a Boot Entry and Setting the Boot Order on HP Integrity Server Blades

Head Node	All Other Nodes
Add a boot entry: 1. From the **Boot Menu** screen, which is displayed during the power on of the node, select the **Boot Configuration Menu**. 2. Select **Add Boot Entry**. 3. Select **Removable Media Boot** as the network boot choice. 4. Enter the string **DVD boot** as the boot option description. 5. Press the **Enter** key twice for no db-profile and load options. 6. If prompted, save the entry to NVRAM.	Add a boot entry: 1. From the **Boot Menu** screen, which is displayed during the power on of the node, select the **Boot Configuration Menu**. 2. Select **Add Boot Entry**. 3. Select **Load File [Core LAN Port 1]** as the network boot choice, which is a Gigabit Ethernet (GigE) port. 4. Enter the string **Netboot** as the boot option description. This entry is required and must be set to the string `Netboot` (with a capital letter `N`). 5. Press the **Enter** key twice for no db-profile and load options. 6. If prompted, save the entry to NVRAM.
Set the boot order: 1. From the **Boot Configuration** menu, select the **Edit OS Boot Order** option. 2. Use the navigation instructions and the up arrow on the screen to move the `DVD boot` entry you just defined to the top of the boot order. 3. Press the **Enter** key to select the position. 4. Press the **x** key to return to the **Boot Configuration** menu.	Set the boot order: 1. From the **Boot Configuration** menu, select the **Edit OS Boot Order** option. 2. Use the navigation instructions and the up arrow on the screen to move the `Netboot` entry you just defined to the top of the boot order. 3. Press the **Enter** key to select the position. 4. Press the **x** key to return to the **Boot Configuration** menu.

For more information about how to work with these menus, see the documentation that came with the HP Integrity server blade.

14. Perform this step on all nodes, including the head node. From the **Boot Configuration** menu, select the **Console Configuration** option, and do the following:
    a. Select **Serial Acpi(HWP0002,PNP0A03,0)/Pci(1|2) Vt100+ 9600** as the primary console interface.
    b. Press the letter **b** key repeatedly until the baud rate is set to 115200.
    c. Press the **Esc** key or press the **x** key as many times as necessary to return to the **Boot Menu**.
    d. If prompted, save the entry to NVRAM.
    e. If prompted, reset the system.

15. Turn off power to the node:
    a. Press **Ctrl+b** to exit the console mode.
    b. Return to the **Command Menu**:

       ```
       MP> CM
       ```

    c. Turn off power to the node:

       ```
       MP:CM> pc -off -nc
       ```

16. Use the **RB** command to reset the BMC.
17. Press **Ctrl+b** to exit the console mode and press the **x** key to exit.

7 Configuring the HP XC System

The information in this chapter parallels the information in Chapter 3 in the *HP XC System Software Installation Guide*. To configure your system, you use a combination of the instructions in this HowTo and the instructions in the *HP XC System Software Installation Guide*. Perform the configuration tasks in the order shown in this chapter.

7.1 Task 1: Install Patches or RPM Updates

Open the *HP XC System Software Installation Guide* to Chapter 3, and follow the instructions in Task 4 to install patches and RPM updates.

7.2 Task 2: Refer To the Installation Guide For System Configuration Tasks

In Chapter 3 of the *HP XC System Software Installation Guide* start with Task 6 and complete all remaining system configuration tasks, which describe how to:

- Set up the system environment
- Run the `cluster_config` utility to configure system services and create the golden image
- Run the `startsys` command to start all nodes and propagate the golden image to all nodes
- Perform other optional and required configuration tasks

Follow the configuration tasks exactly as documented in the *HP XC System Software Installation Guide* with the following exceptions:

- Configuring InfiniBand switch controller cards (Section 7.2.1)
- Starting all nodes and propagating the golden image (Section 7.2.2)

7.2.1 Using Specific IP Addresses to Configure InfiniBand Interconnect Switch Controller Cards

Appendix D in the *HP XC System Software Installation Guide* describes how to configure InfiniBand switch controller cards.

Table D-4 in Appendix D provides a list of the switch names and associated IP addresses to use during the switch controller card configuration procedure. However, because the IP address base differs if the hardware configuration contains server blades and enclosures, you must use the IP addresses shown here in Table 7-1.

Table 7-1 InfiniBand Switch Controller Card Naming Conventions and IP Addresses

Switch Order	Switch Name	IP Address
First switch	`IR0N00`	`172.31.33.1`
Second switch	`IR0N01`	`172.31.33.2`
Third switch	`IR0N02`	`172.31.33.3`
Last switch	`IR0N0`*n*	`172.31.33.`*n*

7.2.2 Running the startsys Command With Specific Options To Start the System and Propagate the Golden Image

Chapter 3 in the *HP XC System Software Installation Guide* describes how to use the `startsys` command to start all nodes and propagate the golden image to all nodes. When you get to that point in the configuration task list, use the following `startsys` options when the hardware configuration contains more than 50 nodes:

1. Propagate the golden image to all nodes:

```
# startsys --image_only --flame_sync_wait=480 --power_control_wait=90 \
    --image_timeout=90
```

2. Boot all nodes when the imaging process is complete:

```
# startsys --power_control_wait=90 --boot_group_delay=45 \
     --max_at_once=50
```

For more information about `startsys` options, see *startsys*(8).

7.3 Task 3: Verify Success

Go to Chapter 4 in the *HP XC System Software Installation Guide* and perform all system verification tasks, including running the Operation Verification Program (OVP) to test successful operation of the HP XC software and hardware components.

7.4 You Are Done

You have completed installing and configuring your HP XC system with HP server blades and enclosures when the OVP reports successful results for all tests.

8 Troubleshooting

This chapter contains troubleshooting information to help you resolve problems you might encounter.

8.1 One or More Ports Do not Communicate Properly on a Gigabit Ethernet Switch

If your hardware configuration contains Gigabit Ethernet switches, verify the following if one of more ports do not communicate properly:

- Switch virtual LANs (VLANs)

  Most managed switches support VLANs. Verify the switch configuration to make sure that all ports you are trying to use are on the same VLAN. Do the same for blade interconnects as well. Check the settings of any relevant switches in any affected server blade enclosures for VLANs.

- Switch management ports

  Some larger ProCurve switches might have a port dedicated for management. This port might not communicate normally with the other ports on the switch until it is configured to do so. Most often, this port is port 1 of slot 1. If you have problems with this port, try using a different port or read the switch documentation for instructions to reconfigure the port or add it to the default VLAN.

8.2 lsadmin limrestart Command Fails

Task 17 in Chapter 3 of the *HP XC System Software Installation Guide* describes LSF postconfiguration tasks. It is possible that the `lsadmin limrestart` command will fail if the LSF control node was assigned to the wrong node name. If the command fails, messages similar to the following are displayed:

```
[root@blc2n1 ~]# lsadmin limrestart
Checking configuration files ...
There are fatal errors.
Do you want to see the detailed messages? [y/n] y
Checking configuration files ...
Platform LSF 6.2 for SLURM, May 15 2006
Copyright 1992-2005 Platform Computing Corporation
Reading configuration from /opt/hptc/lsf/top/conf/lsf.conf
Dec 20 21:00:38 2006 11220 5 6.2 /opt/hptc/lsf/top/6.2/linux2.6-glibc2.3-x86_64-
slurm/etc/lim -C
Dec 20 21:00:38 2006 11220 7 6.2 setMyClusterName: searching cluster files ...
Dec 20 21:00:38 2006 11220 7 6.2 setMyClusterName: Local host blc2n1 not defined
 in cluster file /opt/hptc/lsf/top/conf/lsf.cluster.hptclsf
Dec 20 21:00:38 2006 11220 3 6.2 setMyClusterName(): unable to find the cluster
file containing local host blc2n1
Dec 20 21:00:38 2006 11220 3 6.2 setMyClusterName: Above fatal error(s) found.
------------------------------------------------------
There are fatal errors.
```

To correct this problem, enter the following commands on the head node where *`control_nodename`* is the name of the node that is the LSF control node:

```
# controllsf stop
# controllsf set primary control_nodename
# controllsf start
```

A Configuration Examples

This appendix contains illustrations and descriptions of fully cabled HP XC systems based on interconnect type and server blade height.

The connections are color-coded, so consider viewing the PDF file online or printing this appendix on a color printer to take advantage of the color coding.

A.1 Gigabit Ethernet Interconnect With Half-Height Server Blades

In Figure A-1, only the half-height server blades and the non-blade server nodes have connections to the external network. Because the two NICs of the half-height server blade are already in use, an Ethernet card was added to mezzanine bay 1 to allow for the external network connection on the server blade. An interconnect module was added to bay 3.

On the non-blade server nodes, PCI Ethernet NICs are used for the external network connection. Available ports vary based on hardware model. See the *HP XC Hardware Preparation Guide* for more information about port assignments.

Figure A-1 Gigabit Ethernet Interconnect With Half-Height Server Blades

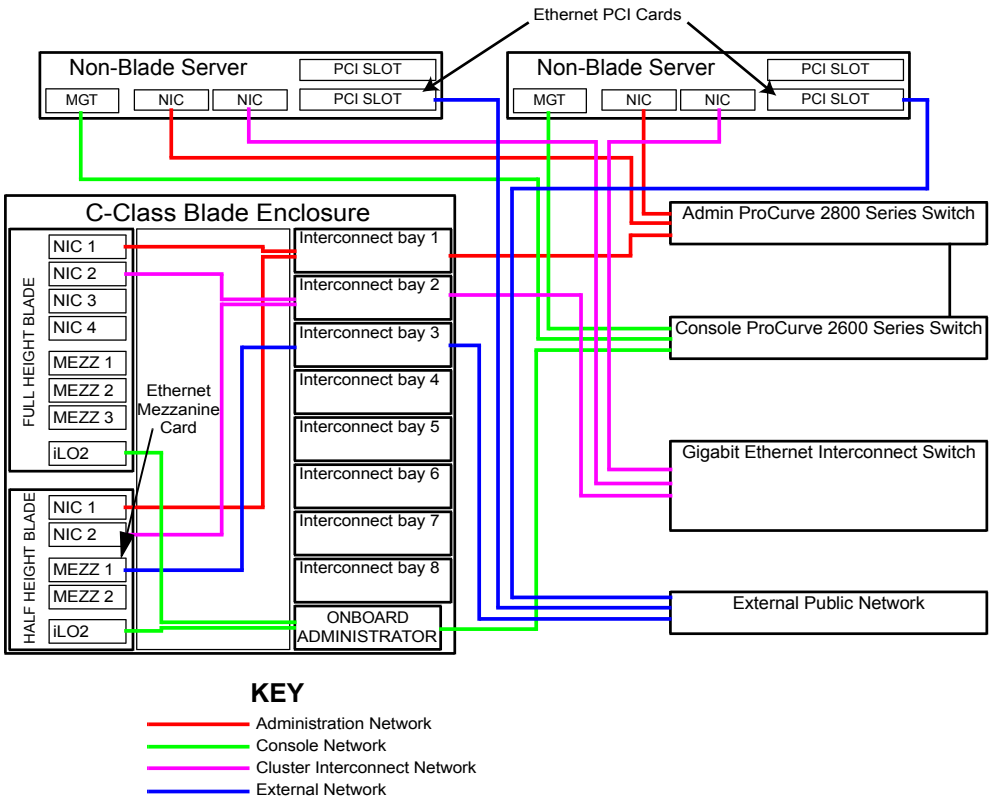

A.2 InfiniBand Interconnect With Full-Height Server Blades

In the configuration shown in Figure A-2, connections to the external network were required only on the non-blade server nodes and the full-height server blades. On those server blades, NIC3 is used for the connection to the external network. A VLAN is used to separate the external network traffic from the administration network traffic on the switch in bay 1 to save the expense of an additional Ethernet interconnect module in bay 2.

On the non-blade server nodes, the built-in NICs were used for the external network connection. Available ports vary based on hardware model. See the *HP XC Hardware Preparation Guide* for more information about port assignments.

Figure A-2 InfiniBand Interconnect With Full-Height Server Blades

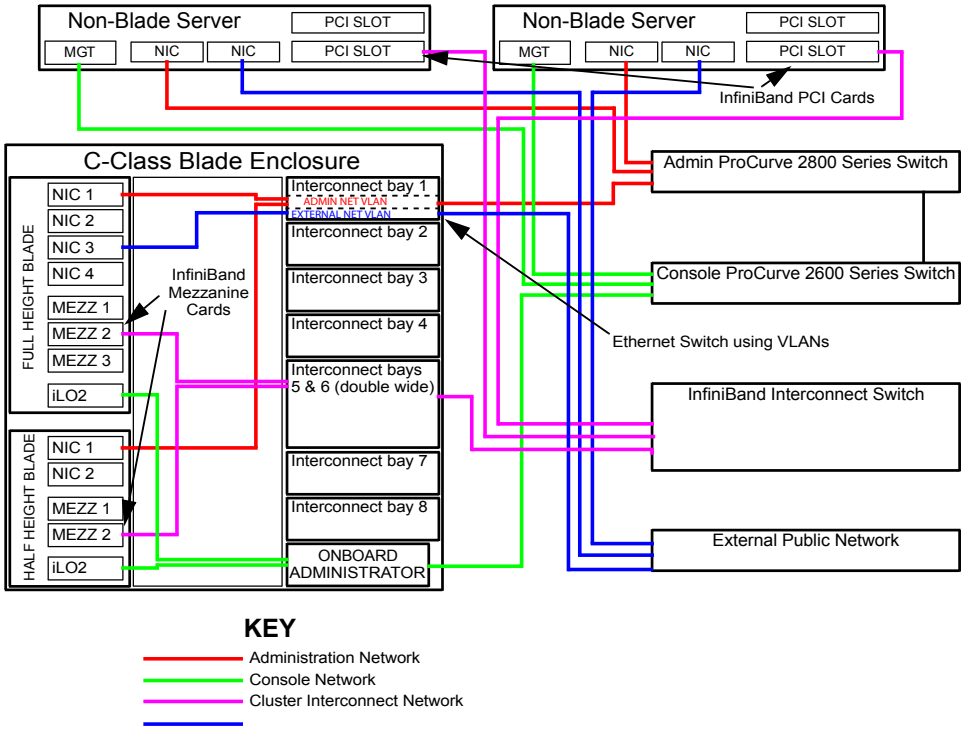

A.3 InfiniBand Interconnect With Mixed Height Server Blades

The configuration shown in Figure A-3 is similar to the configuration shown in Figure A-1 (page 43). The only exception is that in this configuration, the half-height server blades require external connections as well. Because half-height blades have two NICs, you must use NIC2 for the connection to the external network. This also means that an interconnect module is required in bay 2.

On the non-blade server nodes, the built-in NICs are used for the external network connection. Available ports vary based on hardware model. See the *HP XC Hardware Preparation Guide* for more information about port assignments.

Figure A-3 InfiniBand Interconnect With Mixed Height Server Blades

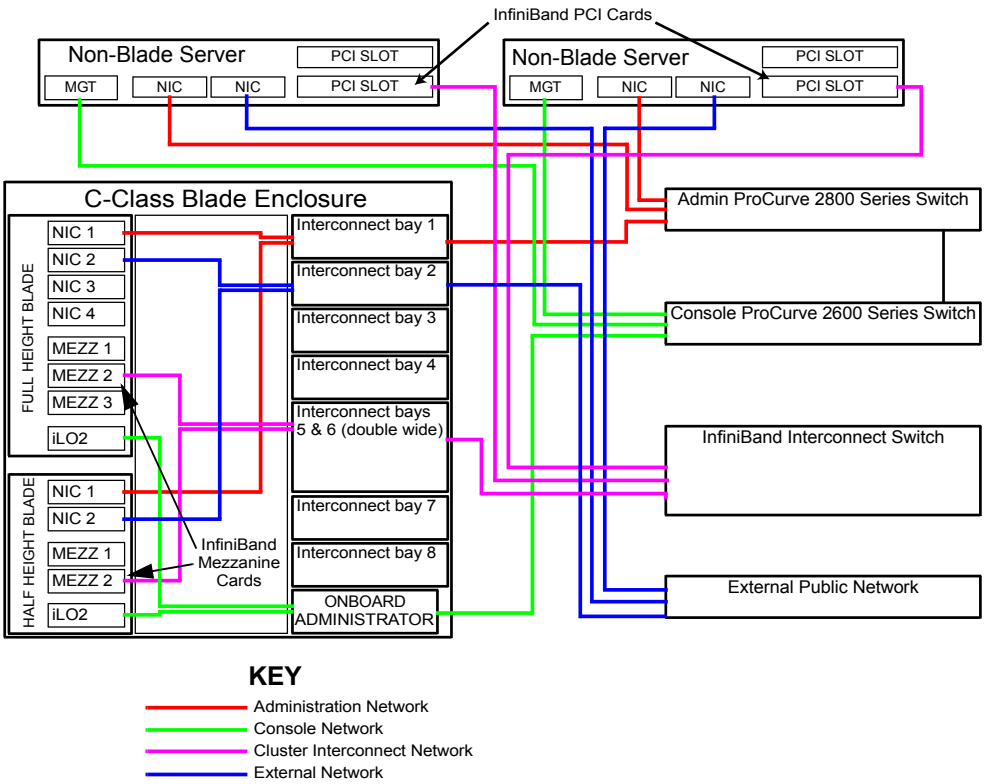

Index

A

administration network
  activating, 29
  as interconnect, 19
  defined, 15

B

base_addrV2.ini file, 28
boot order, 34

C

checklist, 13
cluster_config utility, 39
cluster_prep command, 28
communication problems with GigE switch ports, 41
console network
  defined, 16

D

discover command
  enclosures, 30
  nodes, 30
  switches, 29
discovery process, 27
documentation, 9

E

enclosure
  defined, 11
  discovering, 30
  setup guidelines, 11
/etc/hosts file, 33
external network
  creating VLANs, 19
  defined, 19
  NIC use, 19

G

Gigabit Ethernet interconnect, 17
Gigabit Ethernet switch port, 41
golden image
  propagating to all nodes, 39

H

hardware component
  preparing, 33
hardware components
  defined, 10
hardware configuration
  supported, 10
hardware models
  supported, 10
head node
  installation, 23
  naming convention, 27
HP ProLiant BL460c, 10
HP ProLiant BL465c, 10
HP ProLiant BL480c, 10
HP ProLiant BL685c, 10

I

iLO2
  defined, 11
  features, 11
  setting the password, 34
  virtual console and media features, 23
InfiniBand interconnect, 18
insight display
  defined, 11
installation
  head node, 23
installation procedure, 23
Integrated Lights-Out 2 (*see* iLO2)
interconnect bay port mapping, 12
interconnect module, 12
interconnect network
  defined, 17
  Gigabit Ethernet, 17
  InfiniBand, 18
  running on administration network, 19
IP address base
  base_addrV2.ini file, 28
  changing, 28

L

lsadmin limrestart, 41
LSF
  lsadmin restart failure, 41

M

management processor (*see* MP)
mezzanine cards, 12
MP
  defined, 12
  features, 12

N

network
  administration, 15
  console, 16
  external, 19
  interconnect, 17
network cabling, 15
network configuration, 15
node
  discovering, 30
node name prefix
  defining, 28
node naming, 27

O

onboard administrator
  connecting to, 23
  defined, 11
  IP address, 33
  putting on an active network, 23
  setting the password, 30
  static IP address for, 24
OVP command, 40

P

password
  iLO2, 34
  MP, 37
  onboard administrator, 30
public network (*see* external network)

R

real enclosure
  defined, 27

S

server blade
  boot order, 34
  defined, 9
  preparing HP Integrity nodes, 36
  preparing HP ProLiant nodes, 33
smart array card, 35
startsys command, 39
system configuration, 39

T

telnet access, 34
troubleshooting, 41

V

virtual console and media, 23
virtual enclosure
  defined, 27
virtual local area network (*see* VLAN)
VLAN
  creating, 22
  defined, 19

W

Web site
  documentation, 9